

Express Mail No. EL420322840US

IBM DOCKET: RO999-104
WHE DOCKET: IBM-112

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: FAIR ELEVATOR SCHEDULING ALGORITHM FOR DIRECT
ACCESS STORAGE DEVICE

APPLICANTS: Troy David Armstrong and Michael Steven Faunce

ASSIGNEE: International Business Machines Corporation

Wood, Herron & Evans, L.L.P.
2700 Carew Tower
Cincinnati, Ohio 45202
513-241-2324

SPECIFICATION

652021-1225400

**FAIR ELEVATOR SCHEDULING ALGORITHM FOR
DIRECT ACCESS STORAGE DEVICE**

5 **Field of the Invention**

The invention is generally related to computer access of a shared storage device, and in particular, to scheduling algorithms that arbitrate access between multiple requesters of a shared storage device.

10 **Background of the Invention**

Access to data stored in a storage device can significantly impact the performance of a computer system incorporating such a device. In particular, in high performance or multi-user computer systems, often a relatively large direct access storage device (DASD), e.g., utilizing arrays of magnetic, optical or magneto-optical disk drives and the like, may need to be accessed on a frequent basis by multiple users and/or multiple computer software tasks. DASD's are typically much slower than the solid state memories of such computer systems. As such, the speed at which data can be transferred to or from a storage device can play a large part in the overall performance of a system.

Different users and/or different tasks that require access to a DASD (collectively referred to herein as "requesters") often operate concurrently with one another, and may attempt to access the DASD at the same time. As such, many computer systems employ scheduling algorithms to arbitrate access requests from multiple requesters in a computer system. However, in many instances, it is difficult for a scheduling algorithm to adequately balance the competing interests relating to the performance of the software executing on the computer system, and the performance of the DASD. Whenever these interests are not optimally balanced, overall system performance suffers.

Specifically, DASD performance is often directly impacted by the "seek time" of a DASD -- that is, the time that it takes for a drive head in the DASD to reach the position on a storage medium of the information with which a particular request is associated. One scheduling algorithm that attempts to minimize seek times is referred to as an "elevator" algorithm, which orders requests based upon the relative positions of the data associated with the requests on the DASD. The ordering of the requests is selected so that a drive head can be moved back and forth between the two end positions of the DASD, with any I/O requests encountered along the way issued to the device. The seek time of the drive head on the DASD is thus minimized, thereby maximizing the performance of the DASD.

One potential problem with the elevator algorithm, however, is that if one requester issues a large amount of requests to a relatively narrow range of positions on a DASD, the requester may occupy a large amount of the bandwidth of the DASD. Doing so can effectively restrict access to the DASD for other requesters, thereby stalling the progress of those requesters and adversely impacting system performance.

In other computer systems, a different scheduling algorithm, often referred to as a "fair" algorithm, attempts to schedule requests in a round robin fashion according to the various identities of the requesters associated with the requests such that each requester in a system is able to use the DASD "fairly", and not unduly stall other requesters attempting to concurrently access the DASD. Multitasking performance is often optimized because multiple tasks are allowed to proceed with reduced delays in accessing a shared DASD. However, given that different requesters will often access data stored at different positions on a DASD, the fair algorithm can increase seek times, and thus degrade the performance of the DASD. As such, the fair algorithm is often an optimal solution for multitasking environments and when arbitrating between multiple requesters, but is not as efficient for DASD throughput.

While each of the elevator and fair scheduling algorithms can provide adequate performance in some situations, often the competing interests of ensuring adequate progress of multiple requesters, and minimizing seek times in a DASD, can limit the overall performance of a computer system. Therefore, a significant need

00456211-120799

continues to exist for a manner of scheduling requests from multiple requesters in a computer system to maximize system performance.

2.40 2.41 2.42 2.43 2.44 2.45 2.46 2.47 2.48 2.49 2.50 2.51 2.52 2.53 2.54 2.55 2.56 2.57 2.58 2.59 2.60 2.61 2.62 2.63 2.64 2.65 2.66 2.67 2.68 2.69 2.70 2.71 2.72 2.73 2.74 2.75 2.76 2.77 2.78 2.79 2.80 2.81 2.82 2.83 2.84 2.85 2.86 2.87 2.88 2.89 2.90 2.91 2.92 2.93 2.94 2.95 2.96 2.97 2.98 2.99 3.00 3.01 3.02 3.03 3.04 3.05 3.06 3.07 3.08 3.09 3.10 3.11 3.12 3.13 3.14 3.15 3.16 3.17 3.18 3.19 3.20 3.21 3.22 3.23 3.24 3.25 3.26 3.27 3.28 3.29 3.30 3.31 3.32 3.33 3.34 3.35 3.36 3.37 3.38 3.39 3.40 3.41 3.42 3.43 3.44 3.45 3.46 3.47 3.48 3.49 3.50 3.51 3.52 3.53 3.54 3.55 3.56 3.57 3.58 3.59 3.60 3.61 3.62 3.63 3.64 3.65 3.66 3.67 3.68 3.69 3.70 3.71 3.72 3.73 3.74 3.75 3.76 3.77 3.78 3.79 3.80 3.81 3.82 3.83 3.84 3.85 3.86 3.87 3.88 3.89 3.90 3.91 3.92 3.93 3.94 3.95 3.96 3.97 3.98 3.99 4.00 4.01 4.02 4.03 4.04 4.05 4.06 4.07 4.08 4.09 4.10 4.11 4.12 4.13 4.14 4.15 4.16 4.17 4.18 4.19 4.20 4.21 4.22 4.23 4.24 4.25 4.26 4.27 4.28 4.29 4.30 4.31 4.32 4.33 4.34 4.35 4.36 4.37 4.38 4.39 4.40 4.41 4.42 4.43 4.44 4.45 4.46 4.47 4.48 4.49 4.50 4.51 4.52 4.53 4.54 4.55 4.56 4.57 4.58 4.59 4.60 4.61 4.62 4.63 4.64 4.65 4.66 4.67 4.68 4.69 4.70 4.71 4.72 4.73 4.74 4.75 4.76 4.77 4.78 4.79 4.80 4.81 4.82 4.83 4.84 4.85 4.86 4.87 4.88 4.89 4.90 4.91 4.92 4.93 4.94 4.95 4.96 4.97 4.98 4.99 5.00 5.01 5.02 5.03 5.04 5.05 5.06 5.07 5.08 5.09 5.10 5.11 5.12 5.13 5.14 5.15 5.16 5.17 5.18 5.19 5.20 5.21 5.22 5.23 5.24 5.25 5.26 5.27 5.28 5.29 5.30 5.31 5.32 5.33 5.34 5.35 5.36 5.37 5.38 5.39 5.40 5.41 5.42 5.43 5.44 5.45 5.46 5.47 5.48 5.49 5.50 5.51 5.52 5.53 5.54 5.55 5.56 5.57 5.58 5.59 5.60 5.61 5.62 5.63 5.64 5.65 5.66 5.67 5.68 5.69 5.70 5.71 5.72 5.73 5.74 5.75 5.76 5.77 5.78 5.79 5.80 5.81 5.82 5.83 5.84 5.85 5.86 5.87 5.88 5.89 5.90 5.91 5.92 5.93 5.94 5.95 5.96 5.97 5.98 5.99 6.00 6.01 6.02 6.03 6.04 6.05 6.06 6.07 6.08 6.09 6.10 6.11 6.12 6.13 6.14 6.15 6.16 6.17 6.18 6.19 6.20 6.21 6.22 6.23 6.24 6.25 6.26 6.27 6.28 6.29 6.30 6.31 6.32 6.33 6.34 6.35 6.36 6.37 6.38 6.39 6.40 6.41 6.42 6.43 6.44 6.45 6.46 6.47 6.48 6.49 6.50 6.51 6.52 6.53 6.54 6.55 6.56 6.57 6.58 6.59 6.60 6.61 6.62 6.63 6.64 6.65 6.66 6.67 6.68 6.69 6.70 6.71 6.72 6.73 6.74 6.75 6.76 6.77 6.78 6.79 6.80 6.81 6.82 6.83 6.84 6.85 6.86 6.87 6.88 6.89 6.90 6.91 6.92 6.93 6.94 6.95 6.96 6.97 6.98 6.99 7.00 7.01 7.02 7.03 7.04 7.05 7.06 7.07 7.08 7.09 7.10 7.11 7.12 7.13 7.14 7.15 7.16 7.17 7.18 7.19 7.20 7.21 7.22 7.23 7.24 7.25 7.26 7.27 7.28 7.29 7.30 7.31 7.32 7.33 7.34 7.35 7.36 7.37 7.38 7.39 7.40 7.41 7.42 7.43 7.44 7.45 7.46 7.47 7.48 7.49 7.50 7.51 7.52 7.53 7.54 7.55 7.56 7.57 7.58 7.59 7.60 7.61 7.62 7.63 7.64 7.65 7.66 7.67 7.68 7.69 7.70 7.71 7.72 7.73 7.74 7.75 7.76 7.77 7.78 7.79 7.80 7.81 7.82 7.83 7.84 7.85 7.86 7.87 7.88 7.89 7.90 7.91 7.92 7.93 7.94 7.95 7.96 7.97 7.98 7.99 8.00 8.01 8.02 8.03 8.04 8.05 8.06 8.07 8.08 8.09 8.10 8.11 8.12 8.13 8.14 8.15 8.16 8.17 8.18 8.19 8.20 8.21 8.22 8.23 8.24 8.25 8.26 8.27 8.28 8.29 8.30 8.31 8.32 8.33 8.34 8.35 8.36 8.37 8.38 8.39 8.40 8.41 8.42 8.43 8.44 8.45 8.46 8.47 8.48 8.49 8.50 8.51 8.52 8.53 8.54 8.55 8.56 8.57 8.58 8.59 8.60 8.61 8.62 8.63 8.64 8.65 8.66 8.67 8.68 8.69 8.70 8.71 8.72 8.73 8.74 8.75 8.76 8.77 8.78 8.79 8.80 8.81 8.82 8.83 8.84 8.85 8.86 8.87 8.88 8.89 8.90 8.91 8.92 8.93 8.94 8.95 8.96 8.97 8.98 8.99 9.00 9.01 9.02 9.03 9.04 9.05 9.06 9.07 9.08 9.09 9.10 9.11 9.12 9.13 9.14 9.15 9.16 9.17 9.18 9.19 9.20 9.21 9.22 9.23 9.24 9.25 9.26 9.27 9.28 9.29 9.30 9.31 9.32 9.33 9.34 9.35 9.36 9.37 9.38 9.39 9.40 9.41 9.42 9.43 9.44 9.45 9.46 9.47 9.48 9.49 9.50 9.51 9.52 9.53 9.54 9.55 9.56 9.57 9.58 9.59 9.60 9.61 9.62 9.63 9.64 9.65 9.66 9.67 9.68 9.69 9.70 9.71 9.72 9.73 9.74 9.75 9.76 9.77 9.78 9.79 9.80 9.81 9.82 9.83 9.84 9.85 9.86 9.87 9.88 9.89 9.90 9.91 9.92 9.93 9.94 9.95 9.96 9.97 9.98 9.99 10.00 10.01 10.02 10.03 10.04 10.05 10.06 10.07 10.08 10.09 10.10 10.11 10.12 10.13 10.14 10.15 10.16 10.17 10.18 10.19 10.20 10.21 10.22 10.23 10.24 10.25 10.26 10.27 10.28 10.29 10.30 10.31 10.32 10.33 10.34 10.35 10.36 10.37 10.38 10.39 10.40 10.41 10.42 10.43 10.44 10.45 10.46 10.47 10.

Summary of the Invention

The invention addresses these and other problems associated with the prior art by providing an apparatus, program product and method of processing access requests for a direct access storage device that utilize a "fair elevator" algorithm to schedule
5 access requests from a plurality of requesters desiring access to a DASD. In particular, a fair elevator algorithm consistent with the invention arbitrates requests by taking into account both the requesters with which various requests are associated, as well as the relative positions of the data to be accessed on the DASD. By sorting access requests based upon requester identity and DASD position, multitasking
10 performance and DASD throughput are improved in a balanced manner, thus improving overall system performance.

Consistent with one aspect of the invention, a fair elevator algorithm sorts access requests in two stages. In a first stage, access requests are sorted based upon the requesters associated therewith to generate a first ordered set of access requests.
15 Subsequently, at least a portion of the access requests in the first ordered set of access requests are sorted based upon the positions on the DASD associated therewith. Sorting the first ordered set of access requests creates a second ordered set of access requests, which may be issued in sequence to a DASD to thereby incorporate both requester identity and DASD position into the arbitration of access requests.
20 Consistent with another aspect in the invention, the first and second ordered sets of access requests are stored in queues, with the first queue storing incoming access requests sorted based upon requester identity, and with the second queue receiving access requests from the first queue and sorted based upon the DASD positions associated therewith.

25 Among other advantages, in a multiprocessing or multitasking environment, issuing access requests from multiple tasks may be used to keep multiple processors as active as possible, whereby sorting access requests based upon requester identity maximizes multitasking performance when accessing a DASD. Moreover, grouping access requests that fall close together on a DASD minimizes seek time, and thus
30 improves the response time of individual requests. Moreover, as will become more apparent below, in some embodiments in the invention, it may be desirable to utilize a

5

10

Brief Description of the Drawings

FIGURE 1 is a block diagram of a multiprocessing computer system consistent with the invention.

5 FIGURE 2 is a block diagram of the primary software components in the computer system of Fig. 1, implementing a fair elevator scheduling algorithm consistent with the invention.

FIGURE 3 is a block diagram of an exemplary request data structure for use with the fair elevator scheduling described herein.

10 FIGURE 4 is a flowchart illustrating a process inbound request routine executed by the scheduling algorithm of Fig. 2.

FIGURE 5 is a flowchart illustrating the program flow of a process outbound request routine executed by the scheduling algorithm of Fig. 2.

FIGURES 6A-6E illustrate the processing of an exemplary set of access requests using the scheduling algorithm of Fig. 2.

15 FIGURE 7 is a graph of request issue order versus DASD position for the set of access requests shown in Figs. 6A-6E.

604021-1-123456789

Detailed Description

Turning now to the Drawings, wherein like numbers denote like parts throughout the several views, Fig. 1 illustrates the general configuration of an exemplary computer system 10, or apparatus, suitable for implementing a fair elevator scheduling algorithm consistent with the invention. Computer system 10 generically represents, for example, any of a number of multi-user and/or multi-tasking computers such as a network server, a midrange computer, a mainframe computer, etc. However, it should be appreciated that the invention may be implemented in other data processing systems, e.g., in stand-alone or single-user computer systems such as workstations, desktop computers, portable computers, and the like, or in other programmable electronic devices (e.g., incorporating embedded controllers and the like). For example, one suitable implementation of computer system 10 is in a midrange computer such as the AS/400 computer available from International Business Machines Corporation. Computer system 10 will hereinafter also be referred to as an "apparatus" or "computer", although it should be appreciated the term may also include other suitable electronic devices consistent with the invention.

Computer system 10 generally includes one or more system processors 12 coupled to a memory subsystem including main storage 14 and an external cache system 16. Furthermore, main storage 14 is coupled to a number of types of external devices via a system input/output (I/O) bus 18 and a plurality of interface devices, e.g., an input/output adaptor 20, a workstation controller 22 and a storage controller 24, which respectively provide external access to one or more external networks 26, one or more workstations 28, and/or one or more storage devices such as a direct access storage device (DASD) 30.

Computer system 10 operates under the control of an operating system 40 (shown resident in main storage 14), and executes or otherwise relies upon various computer software applications, components, programs, objects, modules, data structures, etc. Moreover, various applications, components, programs, objects, modules, etc. may also execute on one or more processors in another computer coupled to computer system 10 via a network 26, e.g., in a distributed or client-server

computing environment, whereby the processing required to implement the functions of a computer program may be allocated to multiple computers over a network.

In general, the routines executed to implement the embodiments of the invention, whether implemented as part of an operating system or a specific application, component, program, object, module or sequence of instructions will be referred to herein as "computer programs", or simply "programs". The computer programs typically comprise one or more instructions that are resident at various times in various memory and storage devices in a computer, and that, when read and executed by one or more processors in a computer, cause that computer to perform the operations necessary to execute steps or elements embodying the various aspects of the invention. Moreover, while the invention has and hereinafter will be described in the context of fully functioning computers and computer systems, those skilled in the art will appreciate that the various embodiments of the invention are capable of being distributed as a program product in a variety of forms, and that the invention applies equally regardless of the particular type of signal bearing media used to actually carry out the distribution. Examples of signal bearing media include but are not limited to recordable type media such as volatile and non-volatile memory devices, floppy and other removable disks, hard disk drives, magnetic tape, optical disks (e.g., CD-ROM's, DVD's, etc.), among others, and transmission type media such as digital and analog communication links.

In addition, various programs described hereinafter may be identified based upon the application for which they are implemented in a specific embodiment of the invention. However, it should be appreciated that any particular program nomenclature that follows is used merely for convenience, and thus the invention should not be limited to use solely in any specific application identified and/or implied by such nomenclature.

Those skilled in the art will recognize that the exemplary environment illustrated in Fig. 1 is not intended to limit the present invention. Indeed, those skilled in the art will recognize that other alternative hardware and/or software environments may be used without departing from the scope of the invention.

Fig. 2 illustrates the principal software components suitable for implementing a fair elevator scheduling algorithm consistent with the invention in computer system 10. In particular, operating system 40 is illustrated as interfaced with a DASD 42 to be accessed by a plurality of requesters, here represented by a plurality of tasks 44.

5 Each task 44 is capable of issuing an access request for DASD 42, which is input to a DASD hardware driver 46, sorted based upon a scheduling algorithm 48, and issued in a predetermined sequence to DASD 42 to provide read and/or write access to the DASD.

As used herein, a requester may incorporate any software component capable of accessing a DASD, be it any of multiple computer programs executing on one or more microprocessors, any of one or more computer users attempting to access a shared DASD (each of which may be represented by a separate task in the computer system), any of multiple computers or other devices coupled to a shared DASD or a combination of the above. In general, practically any program code capable of issuing access requests to a DASD in a computer system can operate as a requester in the context of the invention.

A DASD 42 (which may for the purpose of Fig. 2 be considered to logically incorporate both the storage controller 24 and physical device 30 of Fig. 1) may be represented by any form of accessible storage that is capable of being accessed by multiple sources. DASD 42 can incorporate control and drive electronics, and may include an array of multiple physical devices controlled by a central controller. It will be appreciated that, depending upon the storage medium for the DASD, the "position" with which a request is associated may refer to any number of physical locations on the device, including, for example, a logical block address (LBA), an addressable cell, a track number, a head number, a cylinder number, a platter number, and/or a record number, among others.

Likewise, as far as identifying a requester with which an access request is associated, a requester identification may be based upon a specific task identifier, a processor identifier (e.g., so that all requests handled by a given microprocessor are grouped together, irrespective of which task a particular request is associated), a

which identifies whether a request is associated with a read or write operation. Field 76 stores a DASD position associated with the location of the data on the DASD to be accessed by the request. Field 78 stores an operation length indicating the amount of data to access on the DASD starting at the DASD position stored in field 76. In addition, field 80 stores a pointer to a data block 82 that stores the actual information to be written to the DASD in a write-type access request or where the information is to be stored for a read access request. It will be appreciated, however, that request 70 may be implemented utilizing any number of alternate data structures, and the invention is not limited to this particular request data format.

Fig. 4 illustrates a process inbound request routine 90 executed by control logic 58 of Fig. 2. Routine 90 is initiated in response to the receipt of an inbound request from one of the plurality of requester tasks 44. Starting in block 94, routine 90 first determines whether the task queue already has another request with the same requester task ID as the requester task ID associated with the request.

If the task queue does not store any other request for the same requester, control passes to block 96 to insert the request in the task queue based upon the task ID for the request. Essentially, the request is placed in a relative location in the queue based upon its task ID compared to other requests already stored in the queue.

Returning to block 94, if the task queue already holds another request having the same requester task ID, control instead passes to block 98 to insert the request in the task queue immediately after the last request having the same requester task ID to ensure that the relative ordering of requests from the same requester is maintained. Upon completion of either of blocks 96 and 98, routine 90 is complete.

Fig. 5 next illustrates a process outbound requests routine 100 that is invoked when a prior DASD request completes, and a new request needs to be sent to the DASD. Routine 100 typically copies requests stored in task queue 54 over to position queue 60 on an as-needed basis, and resorts the requests based upon DASD positions associated with the requests. Thereafter, requests stored in the position queue are output to the DASD.

In this implementation, routine 100 begins in block 102 by determining whether the position queue is empty. Assuming first that the position queue is

5

10

20

25

30

basis, with the copied requests sorted in an alternating ascending or descending order to minimize seek times in the DASD.

5 The sizes of the queues, as well as the number of requests that are copied at once between the task and position queues, will affect the relative balance of the “fair” and “elevator” aspects of the fair elevator algorithm. Assuming, for example, that position queue 60 processes n requests at a time, a value of $n = 1$ turns the algorithm into a fair algorithm, with only minimal optimization of seek time in the DASD. When $n = \infty$, however, the algorithm operates in a similar manner to a conventional elevator algorithm, where only seek time is optimized. Thus, the
10 selection of an appropriate value of n , as well as the attendant adjustments to the queue lengths and control logic, can be used to operatively “tune” the performance of the algorithm for different environments. Such tuning may be done at design time or during runtime, and may be determined empirically or mathematically.

15 To further explain the operation of the illustrated embodiment of the invention, Table I below illustrates an exemplary set of access requests that may be received by a computer system from a plurality of tasks labeled A-E:

TABLE I: EXEMPLARY SET OF ACCESS REQUESTS

<u>Request No.</u>	<u>Task ID</u>	<u>Operation Type</u>	<u>DASD Position</u>
1	A	Read	101
2	A	Read	102
3	A	Read	103
4	B	Write	80
5	C	Read	930
6	C	Write	513
7	C	Read	933
8	C	Write	514
9	D	Read	333
10	D	Read	881
11	E	Write	731
12	E	Write	732
13	E	Write	733
14	E	Write	734

Assuming, for example, that task queue 54 includes enough entries to store each of the fourteen incoming requests, it will be appreciated that process inbound request routine 90 will order the inbound requests into a first ordered set in the manner shown in Fig. 6A, with the requests sorted by task ID with individual requests having the same ID sorted by the order in which they were received.

Assuming first an ascending sort direction, and assuming at most one request from each task ID is processed each pass, upon pass one of DASD 42, routine 100 pulls the first request for each unique task ID onto position queue 60, and sorts such requests in an ascending order to create a second ordered set of requests.

Subsequently, each of the requests (B, 80), (A, 101), (D, 333), (E, 731), and (C, 930), (where (x, y) refers to a request having a task ID of x and a DASD position of y), is issued in sequence from position queue 60 resulting in the requests for each of the tasks progressing along with little delay and with minimal seek times.

Fig. 6C illustrates the result of a second pass, whereby requests associated with each of tasks A, C, D and E pulled off of task queue 54 and placed on position queue 60, but sorted in a descending manner such that the drive head can return in an opposite direction across the device. Since no pending requests are associated with task B, no task B requests are placed on the position queue.

Figs. 6D and 6E respectively illustrate third and fourth passes of routine 100, wherein the sort direction for position queue 60 is ascending in Fig. 6D and descending in Fig. 6E.

Fig. 7 illustrates in another way the movement of the drive head for DASD 42 when issuing the exemplary requests using the fair elevator scheduling algorithm. The order in which requests are issued, identified by task ID, is shown on one axis of a graph 130, with the other axis showing the relative position on the DASD with which the successively issued requests are associated. It can be seen from the horizontal axis that requests are issued in a relatively fair manner, preventing one task from significantly restricting access to other tasks and requesters in the computer system. On the other hand, as evidenced by the vertical axis, the DASD is capable of sweeping back and forth between end positions with relatively infrequent reversals of direction, thereby minimizing seek times.

Various modifications may be made to the illustrated implementation without departing from the spirit and scope of the invention. For example, rather than copying requests from the task queue to the position queue only when the position queue is empty, the position queue may be filled with new requests when it is determined that the position queue is partially empty, or even on an ongoing basis as new requests are received by the position queue. Moreover, in some implementations, it may be desirable to move multiple requests associated with each task ID from the task queue to the position queue, e.g., to completely fill the position queue during each move operation. In the alternative, only a portion of the position queue may be filled at any given time.

As another additional example, additional factors (e.g., request priority) may also be considered in addition to requester identity and DASD position in a scheduling algorithm consistent with the invention.

Other modifications may be made to the illustrated embodiments without departing from the spirit and scope of the invention. Therefore, the invention lies in the claims hereinafter appended.

00446241-120799